

MULTIPLE VOLTAGE INTEGRATED CIRCUIT AND DESIGN METHOD THEREFOR

CROSS REFERENCE TO RELATED APPLICATION

5

The present invention is related to U.S. Application Serial No. _____
(Attorney Docket No. YOR920030359US1) entitled "METHOD AND PROGRAM
PRODUCT OF LEVEL CONVERTER OPTIMIZATION" to Anthony Correale Jr. et al.,
U.S. Application Serial No. _____ (Attorney Docket No. YOR920030373US1)
10 entitled "SINGLE SUPPLY LEVEL CONVERTER" to Anthony Correale Jr. et al., both
filed coincident herewith and to U.S. Application Serial No. 10/387,728 (Attorney
Docket No. RPS9-2002-0253) entitled "VOLTAGE ISLAND CIRCUIT PLACEMENT"
to Anthony Correale Jr., filed March 13, 2003, all assigned to the assignee of the present
invention.

15

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention is related to integrated circuit (IC) design circuit design and
20 more particularly, to optimizing standard cell design configurations.

Background Description

Semiconductor technology and chip manufacturing advances have resulted in a
25 steady increase of on-chip clock frequencies, the number of transistors on a single chip
and the die size itself, coupled with a corresponding decrease in chip supply voltage and
chip feature size. Generally, all other factors being constant, the power consumed by a
given clocked unit increases linearly with the frequency of switching within it. Thus, not
withstanding the decrease of chip supply voltage, chip power consumption has increased

as well. Both at the chip and system levels, cooling and packaging costs have escalated as a natural result of this increase in chip power. For low end systems (e.g., handhelds, portable and mobile systems), where battery life is crucial, net power consumption reduction is important but, must be achieved without degrading performance below acceptable levels. Consequently, power consumption has been a major design consideration for designing very large scale integrated circuits (VLSI) such as high performance microprocessors. In particular, increasing power requirements run counter to the low end design goal of longer battery life. Since chip power is directly proportion to the square of supply voltage (V_{dd}), reducing supply voltage is one of the most effective ways to reduce the power consumption, both active and standby (leakage) power, which is becoming more and more of a problem as technology features scale into nanometer (nm) dimension range.

While reducing supply voltage is attractive to reduce the power consumption, reducing V_{dd} increases transistor and gate delay. Thus, for a design that is performance constrained, the supply voltage may not be lowered too much and, it is usually determined by the *most* timing critical paths. However, it is often the case that most cells in a chip are timing non-critical. If those timing non-critical cells are properly selected to be on lower supply voltage(s), significant power saving may be achieved without degrading the overall circuit performance.

One approach to reducing power is to use multiple supply voltages each supplying different circuit blocks or voltage islands. Each voltage island runs at its minimum necessary supply voltage. However, multiple supply voltages on the same circuit/chip present numerous problems, especially for deep submicron (DSM) designs, where circuit performance often is dominated by interconnect delays. In particular, logic synthesis is very complicated for multiple supply designs and, placement and routing must be considered together for voltage assignment, level converter insertion and minimization, and for circuit block clustering to simplify power routing of multiple supply lines.

Thus, there is a need for circuit element clustering for minimum power and to simplify power routing of multiple supply lines.

SUMMARY OF THE INVENTION

5 It is a purpose of the invention to improve integrated circuit (IC) chip design;

 It is another purpose of the invention to improve cell placement in multi supply voltage IC chip designs;

 It is yet another purpose of the invention to improve cell placement of first supply voltage cells with cells of other supply voltages in multi supply voltage IC chip designs;

10 It is yet another purpose of the invention to group circuit cells in a multi-supply design close to their respective power supplies;

 It is yet another purpose of the invention to group circuit cells in a multi-supply design to facilitate timing closure;

15 It is yet another purpose of the invention to group circuit cells in a multi-supply design for optimum level converter placement;

 It is yet another purpose of the invention to group circuit cells in a multi-supply design for a minimum number of level converters;

 It is yet another purpose of the invention to group circuit cells in a multi-supply design for efficient level converter placement.

20 The present invention relates to an integrated circuit (IC) design, method and program product for reducing IC design power consumption. The IC is organized in circuit rows. Circuit rows may include a low voltage island powered by a low voltage

(V_{ddl}) supply and a high voltage island powered by a high voltage (V_{ddh}) supply. Circuit elements including cells, latches and macros are placed with high or low voltage islands to minimize IC power while maintaining overall performance. Level converters may be placed with high voltage circuit elements.

5

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

10

Figures 1A – C show different state of the art circuit layouts for multi-supply chips;

Figure 2 shows an example of a generic voltage island structure formed according to a preferred embodiment of the present invention;

15

Figure 3 shows an example of a flowchart of a method of generic voltage island optimization for low power with rapid timing closure according to a preferred embodiment of the present invention;

Figures 4A – B show an example of the steps in the logic aware voltage assignment;

20

Figures 5A – B show an isolated V_{ddl} cell (e.g., width 1 cell) in the middle of a larger V_{ddh} island, optimized by changing such isolated cells back to a V_{ddh} cell;

Figures 6A – F show before and after level converter placement examples, optimized according to a preferred embodiment of the present invention;

Figures 7A – B show an example of a V_{ddl} fanin cone for an iterative level converter optimization;

Figure 8 shows an example of level converter efficiency measurement flow diagram using V_{ddl} fanin cone size to iteratively locate and delete least efficient level converters;

Figures 9A – B show before and after examples of level converter optimization effected with logic replacement;

Figure 10 shows a flow diagram showing an example of the logic replacement;

Figures 11A –B show before and after examples of replacing a buffer and level converter with a single level converter and adjusting placement to meet design objectives;

Figure 12 shows a flow diagram for identifying paired level converters and buffers for optimization.

DESCRIPTION OF PREFERRED EMBODIMENTS

Accordingly, as described hereinbelow, the present invention provides a versatile and generic multi-supply voltage island circuit structure, wherein different supply voltages are assigned at both macro and cell level within the islands. Unless indicated otherwise for simplicity of discussion hereinbelow, logic cell and gate are used interchangeably and each is a sub-circuit of standard cell design. Further, a standard cell design is taken as having the same height, i.e., row height, for most cells. Abutting cells form circuit rows. Also, typical modern application specific integrated circuit (ASIC) and system on a chip (SOC) designs often have many proprietary macros (known in the art as intellectual property (IP) blocks) mixed with standard cells. A voltage island can be a single cell, an IP block or macro or, a continuous region of cells on the same or

adjacent rows that have the same power supply voltage (referred to herein as a high voltage supply or V_{ddh} and a low voltage supply or V_{ddl}). An output or source drives a net connecting one or more inputs or sinks to the source and a low/high voltage net connects a low/high voltage source to low/high voltage sinks. Also, although described herein in terms of two (2) supplies description, this is for example only and not intended as a limitation. A person skilled in the art would readily understand how to extended the 2 supply description to multiple supply voltages.

So, Figures 1A – C show different state of the art multi-supply chips with examples of well known circuit island placement, e.g., as described in D.E. Lackey et al., “Managing power and performance for system-on-chip designs using voltage islands”, in *Proc. International Conference on Computer Aided Design*, pp. 195 – 202, Nov. 2002. In the example 100 of Figure 1A, voltage islands are only allowed at the macro level 102, 104, with no fine-grained voltage assignment for cells 106, 108. For deep submicron (DSM) designs, which have tight performance targets, it may not be possible to switch an entire macro between a normal and a lower supply voltage without incurring an overall circuit performance loss. So it would be more flexible if voltage assignment can be done at cell level to exploit positive slacks. The example of Figure 1B shows a circuit block 110 with cell level voltage assignment, but at the cost of a restricting the layout to alternating or interleaving pairs of high and low supply rows 112, 114. Figure 1C shows another example 120, somewhat unconstrained by the requirement of interleaving entire rows. Instead, in this example each row 122, 124, 126, 128 may have two areas with different voltages (designated H or L), provided each area occupies either the left or right part of the row. Unfortunately, these voltage island patterns or segregated voltage areas unnaturally constrain voltage assignment and/or reduce placement flexibility. Frequently in a typical modern ASIC/SOC design, non-critical regions are interspersed with critical regions in the same circuit row. Typically available such circuit structures are not flexible enough to allow circuit placement or voltage island granularity sufficient to meet stringent delay constraints or, in placing to meet such constraints introduce routing problems.

By contrast, a preferred circuit and chip design method incorporates generic voltage islands with much finer layout granularity. Supply voltage assignment may be done at both macro and gate level, affording designers much more design freedom and providing a much more flexible voltage island layout structure. Further such a preferred embodiment design achieves timing closure on design timing goals during voltage island formation and hastens timing optimization.

Figure 2 shows an example of a generic voltage island structure 130 formed according to a preferred embodiment of the present invention, wherein different voltages are assigned at both macro and cell levels. Preferred voltage assignment affords more freedom in terms of layout style by allowing multiple voltage islands within the same circuit row. Further, such a pattern 130 is achievable with minimum disturbance to an existing placement, i.e., after normal chip design and placement. So, after designing and placing circuits for performance, for example, the design may be modified according to the present invention, selectively replacing higher power (V_{ddh}) circuits (stippled) with lower power (V_{ddl}) circuits (clear) where possible. Since some gap may be needed between adjacent V_{ddl} islands 132 and V_{ddh} islands 134 (depending on the standard cell library), a minimum or maximum allowed cluster size or number of voltage islands may be specified for each circuit row, e.g., 136, based on the particular user or technology specification. See, for example, U.S. Application Serial No. 10/387,728 (Attorney Docket No. RPS9-2002-0253) entitled "VOLTAGE ISLAND CIRCUIT PLACEMENT" to Anthony Correale Jr., filed March 13, 2003, assigned to the assignee of the present invention and incorporated herein by reference. To facilitate power routing, a power grid structure of VDDL 138 and VDDH 140 is co-designed with the voltage island assignment.

Typically, a V_{ddl} source cannot drive a V_{ddh} sink reliably without excessive leakage. Thus, a level converter is needed for a transition from a low voltage net to a high voltage net. Traditional level converters require both supply voltages, V_{ddl} and V_{ddh} , to avoid excessive leakage. Previously, using dual-supply voltage level converters 144,

required that they be placed at the island 132, 134 boundaries for access to both power supplies. However, a single-supply level converter is used such as is described in U.S. Application Serial No. _____ (Attorney Docket No. YOR920030373US1) entitled "SINGLE SUPPLY LEVEL CONVERTER" to Anthony Correale Jr. et al., filed
5 coincident herewith and incorporated herein by reference. Correale Jr. et al. level converters 142 can be placed anywhere in a higher voltage island 134 or logic 146 and so, provide additional placement flexibility. Preferably, a level converters as described hereinbelow is a single supply level converter such as Correale Jr. et al.

Figure 3 shows an example of a flowchart 150 of a method of generic voltage
10 island optimization for low power with rapid timing closure according to a preferred embodiment of the present invention. For deep submicron (DSM) designs, interconnect delay can dominate the transistor delay, thus placement (and even routing) information are used to get an accurate timing estimation.

So, beginning in step 152 an input netlist description and specifications (e.g.,
15 technology files and timing constraints) is provided. In step 154 a timing closure tool with Spice RC delays (e.g., a suitable tool from Synopsis, Inc., or EinsTimer from IBM Corporation) is used to determine the entire circuit/chip timing at the higher supply voltage (V_{ddh}) for a base placement and optimization, i.e., determining global placement and obtaining a good timing estimation. Then, non-critical cells are identified and
20 assigned a lower supply voltage (V_{ddl}). As noted hereinabove, interconnect delay can dominate the gate delay for deep submicron circuits and so, power can be reduced for lightly loaded circuits where power is not needed for driving large interconnect loads. So, the global placement information is used to correctly identify the critical versus non-critical cells, e.g., heavily loaded verses lightly loaded. Then in step 156, a logic aware
25 voltage assignment is performed, assigning the lower supply voltage(s) to less critical circuits, i.e., macro, latch and/or cell. Next, in step 158 level converters are inserted and the results are refined and optimized. A level converter is inserted wherever there is a transition net with a low voltage cell driving a high voltage cell or, where a pass gate data

input to a low voltage cell or circuit element is being driven by a high voltage cell and being controlled by a low voltage cell. In step 160 isolated assignments are removed in a physical aware voltage reassignment step, locating and reverting solo or very small groups of low voltage circuits that are difficult to form into low voltage islands. Since eliminating those isolated low voltage cells may create opportunities to reassign previously assigned high voltage cells to low voltage cells, in step 162 the design is checked for such opportunities. If any are found, returning to step 156 for another pass the design is further optimized, until there is no improvement available in step 162. Finally, in step 164 placement and power routing patterns are effected based on the voltage island assignments to form the final high and low voltage islands. As result, the entire flow can be tightly integrated with a suitable physical synthesis engine 166 such as a routing tool from Cadence Design Systems, e.g., for application of any necessary further timing optimization.

In addition to identifying circuits for separation into voltage islands, supply high and low voltages may similarly be selected to achieve optimum power saving. Further, a preferred voltage assignment method has application to static and incremental timing engines. Every time a macro or cell is changed from a higher voltage cell to a lower voltage cell, or vice verse, the timing (slack) is updated.

Figures 4A – B show an example of the steps in the logic aware voltage assignment step 156 of Figure 3. Essentially, a logic assessment is done for each macro 1560, latch 1562 and cell 1564 to determine which may be replaced with a low voltage equivalent and level converter, if required. For checking combinational logic cells in step 1564, the cells may be sorted according to timing order from timing end point to timing starting point, i.e., from PO to PI or latch input to latch output. In each major step 1560, 1562 and 1564, each circuit element of each group (macro, latch or cell) is checked, essentially according to the steps 1570 – 1576 in Figure 4B to identify low voltage candidates. First in step 1570, the supply to the macro, latch or cell is reduced and one or more level converters are inserted where appropriate, i.e., at transition nets with low

voltage sources driving high voltage sinks. In step 1571 an appropriate incremental timing report is checked for the macro, latch or cell. Then, in step 1572, if the timing specification of the macro, latch or cell is met, it is designated to the low supply voltage. For latches in particular, a latch is designated a low supply latch, if all input pins still
5 have positive slack (i.e., edges arrive at inputs prior to a minimum input set up time) and the output pin slack exceeds a minimum threshold, i.e., for a transitional net the output can accommodate the additional delay for an inserted level converter. Otherwise, in step 1573 it is reverted to the normal, higher supply. In step 1574, if additional macros, latches or cells have not yet been checked; then in step 1575, the next (macro, latch or
10 cell) is selected and returning to step 1570, checking continues. Once, each element of the particular group being checked, i.e., in step 1560, 1562 or 1564, checking proceeds to the next group in 1562 or 1564, respectively, or ends in step 1576. After an initial voltage assignment, the voltage assignment may be further refined, including deleting smaller low voltage supply clusters.

15 The initial voltage assignment is not physically aware, i.e., no consideration is given to cell placement. As shown in the example of Figure 5A, it is possible to assign an isolated V_{ddl} cell 170 (e.g., width 1 cell) in the middle of a larger V_{ddh} island, 172A – B, 174, 176, 178. Since such an isolated placement may make it difficult to form uniform voltage islands, an optimum placement is facilitated by changing each such isolated cell
20 170 back to a V_{ddh} cell 170' as shown in Figure 5B. It should be noted that initial assignment of these isolated V_{ddl} cells may have prohibited considering other V_{ddl} cells as candidates. Thus, a physical aware voltage reassignment is employed to push more cells to V_{ddl} while minimizing the number of level converters and still meeting the physical timing constraints. So, physical adjacency information is used to facilitate the physical
25 aware voltage reassignment and to guide subsequent voltage assignment.

Physical aware voltage reassignment step 160 in Figure 3, basically, includes 2 steps. First, a physical adjacency metric (PAM) is computed for the each V_{ddl} cell. The PAM(k, d) for each particular V_{ddl} cell is, the total size (i.e., width) of V_{ddl} cells within the

neighboring k rows, including the cell itself, and within diameter range d . Then, all V_{ddl} cells with a PAM less than certain threshold are reverted to V_{ddh} cells. Each reversion may present new opportunities for converting some other V_{ddh} cells that had not been selected in the initial voltage assignment, e.g., due to slack constraints. So, in step 162 of Figure 3 logic aware voltage assignment is called again with PAM as an additional metric. Only those cells with PAM larger or equal to the selected threshold may be selected as V_{ddl} cells. Thus, the logic aware assignment step 156 and physical aware reassignment step 162 may be iterated until no further improvement is realized.

In each iteration level converter placement is optimized in step 158 to reduce the total number of level converters, gradually deleting the less efficient level converters. Level converters are necessary for transitions between islands, i.e., at least when a V_{ddl} source is driving a V_{ddh} sink. So, for example, branches to those level converters with a small V_{ddl} fanin may be eliminated (deleting the level converter and returning the prior cell with a V_{ddl} cell) or another level converter efficiency metric may be used to select level converters for deletion. Further, since level converters and buffers essentially have the same function and so, can be substituted for buffers, optimizing level converters, simultaneously optimizes buffers. In particular, for any V_{ddl} output driving multiple V_{ddh} inputs (i.e., inputs to multiple V_{ddh} cells), instead of inserting a level converter for each V_{ddh} input, a single level converter is shared, provided that timing and electrical constraints are still met.

Figures 6A – F show before and after level converter placement examples. In the example of Figure 6A, a V_{ddl} driver 180 is shown driving a transition net with two V_{ddh} receivers 182, 184 aligned in a straight line, where the level converter 186 is at the geometric center of the two receivers 182, 184. However, this placement increases the total wire length because of the detour from the driver 180 to the level converter 186 and, then to the left receiver 182. By contrast, as shown in Figure 6B, an optimized placement places the level converter 186 just in front of the left receiver 182 to minimize the total

net power by maximizing the low voltage net length portion. Thus, power saving may not necessarily decrease the total wire length, but optimizes its apportionment.

Similarly, as shown in the examples of Figures 6C – D, placement can be optimized for a driver 190 driving a transition net with receivers 192, 194, 196, 198 on a two dimensional plane from the driver 190. In this example, the receivers 192, 194, 196, 198 are all located in the first quadrant from the perspective of the driver 190. A common level converter 200 can be shared between V_{ddl} and V_{ddh} interfaces. Preferably, however, the optimum level converter 200 placement is a location to minimize the total wire length; and also, allocates the largest portion of that wire length to the low supply voltage side (i.e., driven by the V_{ddl} driver 190) to minimize switching power, i.e., power expended driving the wire load. Thus, in the example of Figure 6C the level converter 200 is located a minimum power point at (X_{min}, Y_{min}) , where X_{min} and Y_{min} are the minimum x and y coordinates of all receivers 192, 194, 196, 198. Thus selecting the minimum power point avoids any total wire length increase, but may place the level converter 200 closer to the driver 190. Alternatively, in Figure 6D the level converter 200 may be placed at the Manhattan distance from the nearest sink (receiver 194 in this example) to the source (on the 45° dotted line 202 in this example). A weighted geometric center 204 may be determined for all the receivers 192, 194, 196, 198 from a delay neutral drive point from the level converter 200. The weight applied for each receiver 192, 194, 196, 198 is a measure of how close the receiver should be to the driver 190 (e.g., the weight may be measured by the slack at each receiver). Then, a projection is determined from the weighted geometric center 204 to the 45° dotted line 202 is performed to determine the level converter location. The weighted center placement more aggressively pushes the level converter 200 further away from the source 190 to increase the total V_{ddl} wire length and thus reduce V_{ddh} wire length, and as a result, minimize power.

Figures 6E – F show after placement examples, wherein V_{ddh} receivers 210, 212, 214, 216 are located in more than just a single quadrant, e.g., they occupy both the first

and the fourth quadrant. In this example, the level converter 218 is placed at a side drive point (X_{\min} , Y_{drv}), where X_{\min} is the minimum x-coordinate of all receivers, and Y_{drv} is the y-coordinate of the driver 220. Similar drive points can be located for first-second quadrants, second-third quadrants, and third-fourth quadrants. However, if as in the
5 example of Figure 6F, the receivers 230, 232, 234, 236 238 are dispersed in diagonal quadrants (e.g., first-third quadrants, or second-fourth quadrants), the level converter 240 is placed near the driver 242 because it may not be inserted at any other place without increasing the total wire length.

It should be noted that in all of the above examples, if one level converter 186,
10 200, 218, 240 is not enough to drive all the respective V_{ddl} receivers, it may be powered up using any suitable technique, e.g., cloning. Whether the level converter is powered up through cloning or otherwise should be evaluated together with the overall power saving of the placement. In particular, the original assignment of V_{ddl} driver may be reverted to V_{ddh} if the level converter cost is higher than the gain by selecting the driver to be V_{ddl} in
15 the first place. Furthermore, level converter placement as described with reference to Figures 6A – F is done focusing on total power saving, by minimizing the overall capacitance and V_{ddh} cell load capacitance, while maximizing the V_{ddl} cell load capacitance after level converter placement. However, application of the above described level placement may be done guided by any other selected cost function, such as timing
20 and power supply adjacency, i.e., to deliver proper power supplies to level converters. After the level converter is inserted and placed, a Steiner tree is constructed to connect the level converter with the V_{ddh} receivers.

Figures 7A – B show an example of an iterative optimization of level converter placement for a V_{ddl} fanin cone 250 according to a preferred embodiment of the present
25 invention. Generally, a fanin cone for level converter includes all gates that drive nets leading to the gate inputs and, as applied to the level converters, signals originating from V_{ddl} gates without crossing/passing through any V_{ddh} gates. As a rule of thumb, the larger the V_{ddl} fanin cone, the more effective the level converter.

In this example the V_{ddl} fanin cone 250 for level converter 252 includes the 5 gates 254, 256, 258, 260, 262. In this example, the size of each V_{ddl} fanin cone for the level converters 252, 266 and 268 is 5, 1 and 4, respectively. However, since each level converter 252, 266, 268 consumes power and chip area, placement is optimized by deleting inefficient level converters. To the first order, the size of V_{ddl} fanin cone is a rough measure of the efficiency of a particular level converter. So, level converters that are inefficient, i.e., level converters with small fanin cones, are deleted. For example, the level converter 266, which has V_{ddl} fanin cone size of one (i.e., only one buffer 270 driving into it) and so, may not be cost effective with respect to power or area. Further, as shown in Figure 7B after deleting level converter 266 and reverting the single, low voltage input buffer 270 to V_{ddh} buffer 272, the inefficient fanin cone has been eliminated. Also, after deleting level converter 266, the V_{ddl} fanin cones of level converters 252 and 268 are 4 and 4, respectively.

Figure 8 shows an example of level converter efficiency measurement flow diagram 280 using V_{ddl} fanin cone size to iteratively locate and delete least efficient level converters according to a preferred embodiment of the present invention. First, in step 282 the V_{ddl} fanin cone of each level converter is determined. Then, in step 284 level converters with a fanin having a cone size less than or equal to a selected threshold, k , are converted to V_{ddl} cells. Next in step 286 the V_{ddl} fanin cone size for remaining level converters is updated. In step 288 fanin cones are checked to determine whether more inefficiently placed level converters can be removed, i.e., have a fanin cone size below k . If more fanin cones with a size below k remain, then, returning to step 284, remaining such inefficient level converters are removed, one at a time until none are found in step 288 and optimization ends in step 290. Further, a minimum threshold of V_{ddl} fanin cone size k_{min} may be obtained, incrementally, or a total level converter number upper bound may be incrementally increased to gradually reach an optimum placement. So, the bound may be incrementally increased, gradually removing least efficient level converters, i.e., by setting $k = 1$ first, then $k = 2, 3$, and so on until $k = k_{min}$ or until a selected total level converter number requirement is met. It should be noted also that using V_{ddl} fanin cone

size as described herein as a level converter efficiency metric is for example only and not intended as a limitation. Any other measurement metric may be employed to iteratively select and delete less efficient level converters.

Figures 9A – B show before and after examples, 300, 302, respectively, of level converter placement optimization effected with logic replacement, i.e., replacing selected V_{ddh} gates with its V_{ddl} counterparts (possibly using a different size in the library) to reduce the number of level converters. In particular, this is effective for those V_{ddh} gates that have many fanin signals originating with level converters. So for example, in before circuit 300 gate 304 is assigned to V_{ddh} , because it is timing critical due to another input from a V_{ddh} gate 306. The gate 304 receives its four other inputs from gates 308, 310, 312, 314 that are all V_{ddl} cells and so, require insertion of four level converters 316, 318, 320, 322. Thus, in optimized circuit 302, gate 304 is replaced with a functionally equivalent V_{ddl} gate 324 and, typically, a level converter (not shown) is inserted at output 326. In addition, the replacement V_{ddl} gate 324 may be of a different drive strength. However, the number of level converters may be significantly reduced by such replacement.

Figure 10 shows a flow diagram showing an example of the logic replacement step 330 according to a preferred embodiment of the present invention. First, in step 332 a V_{ddh} gate candidate with multiple input level converters is identified. Then, in step 334 the selected V_{ddh} gate is temporarily replaced with its V_{ddl} equivalent. Unnecessary level converters are deleted from the inputs to the replaced gate and, if necessary, a level converter is inserted at the gate output. Then in step 336, the timing constraint is checked to determine if it is still met. Optionally, step 334 may be repeated, trying different V_{ddl} gate sizes and selecting the best result for timing/power. If timing is met in step 336, then the logic replacement with the most power saving is selected in step 338. Otherwise, in step 340 the previous (original) solution is restored. In step 342 the logic is checked to determine if more V_{ddh} candidates remain. If so returning to step 322 the next

V_{ddh} candidate is selected, until in step 342 no candidates remain and so, all candidate V_{ddh} gates with multiple level converters in its inputs are checked.

Figures 11A –B show before and after examples 350, 352, wherein a buffer 354 and level converter 356 are replaced, with a single level converter 358 and placement is adjusted to meet design objectives. As noted hereinabove, since each level converter is itself a buffer, level converters can be substituted for traditional buffers, e.g., as signal relays to break long interconnects and restore/redrive signals, thereby reducing buffers or chains of inverters.

Figure 12 shows a flow diagram 360 for identifying paired level converters and buffers for optimization. First in step 362, a each level converter is identified with at least one buffer immediately before it with fanout 1 (FO1). If such a level converter is identified, then in step 364 the buffer is temporarily removed, and the level converter placement is adjusted as described hereinabove. Then in step 366, the timing specification is checked and, if still met, the buffer is permanently removed. Otherwise, in step 368, the original placement is restored. Then, in step 370 the remaining buffers are checked for more candidates and, if one is found, returning to step 364, that candidate is checked. Otherwise, checking ends in step 372.

A design may be constrained wherein portions may not be modified, e.g., with input/output (I/O) constraints that may not be replaced, for example, with V_{ddl} cells. For example in a microprocessor core design, placing slower V_{ddl} cells at the input logic between primary chip input and the first level latches, as well as at the output logic between the final level latches and the primary chip outputs may be unacceptable. Such constrained logic can be hidden or removed from consideration to avoid changing those cells to V_{ddl} cells. Also, a user may specify a supply voltage for a set or sets of cells or macros. Such constraint information can be passed to voltage assignment with those constrained cells marked as hidden and so, not touched. Also, circuitry related constraints, can be applied during voltage assignment.

Advantageously, the present invention provides a flexible, systematic method for identifying cell candidates and creating optimized voltage islands. Further, such a design is achieved with a fine-grained voltage island and without performance degradation. Additionally, voltage assignment is both logically and physically, honoring both logic and physical adjacencies. Level converters are efficiently optimized for the design.

While the invention has been described in terms of preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.